



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification 5 :	A1	(11) International Publication Number: WO 91/11765
G06F 9/40		(43) International Publication Date: 8 August 1991 (08.08.91)

(21) International Application Number: PCT/US91/00587

(22) International Filing Date: 29 January 1991 (29.01.91)

(30) Priority data:
471,962 29 January 1990 (29.01.90) US

(71) Applicant: T RAPLEX, INC. [US/US]; 100 Trade Center Drive, Suite 301, Champaign, IL 61820-7237 (US).

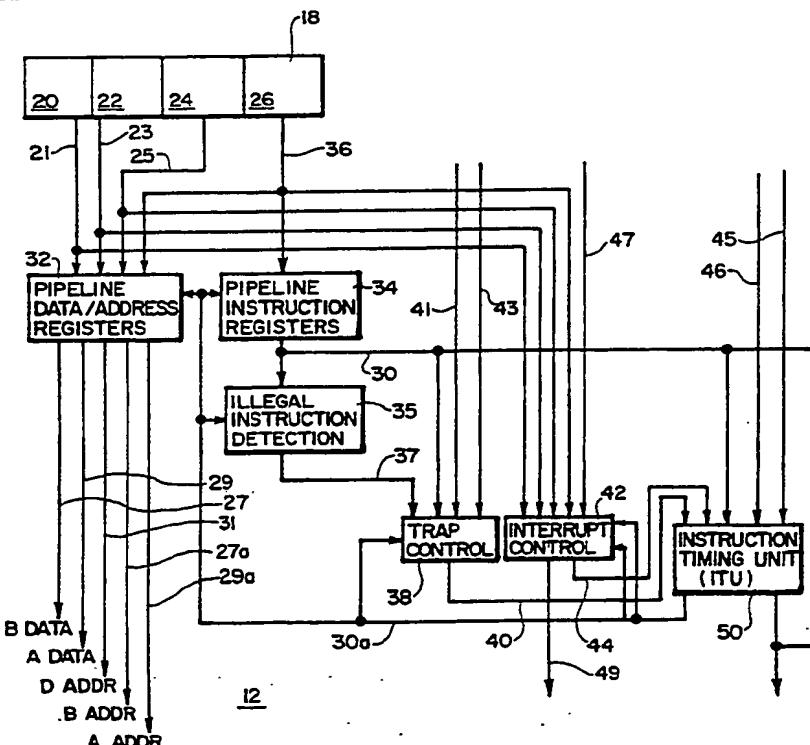
(72) Inventor: GLICKMAN, Jeff, Bret ; 2802 Prairie Meadow Road, Champaign, IL 61821 (US).

(74) Agents: KRIEGER, Daniel, J.; Neuman, Williams, Anderson & Olson, 77 West Washington Street, Chicago, IL 60602 (US) et al.

(81) Designated States: AT (European patent), AU, BE (European patent), CA, CH (European patent), DE (European patent), DK (European patent), ES (European patent), FR (European patent), GB (European patent), GR (European patent), IT (European patent), JP, KP, KR, LU (European patent), NL (European patent), SE (European patent), SU.

Published
With international search report.

(54) Title: ARCHITECTURE FOR MINIMAL INSTRUCTION SET COMPUTING SYSTEM



(57) Abstract

A computing system (14) architecture that uses a minimal instruction set for the functioning of a general purpose computing system as well as applying completely unencoded instructions from memory (16) directly to the hardware is herein described. The present invention additionally uses a flowthrough design to further reduce the hardware complexity to provide a streamlined and extremely efficient architecture (10).

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AT	Austria	ES	Spain	MG	Madagascar
AU	Australia	FI	Finland	ML	Mali
BB	Barbados	FR	France	MN	Mongolia
BE	Belgium	GA	Gabon	MR	Mauritania
BF	Burkina Faso	GB	United Kingdom	MW	Malawi
BC	Bulgaria	GN	Guinea	NL	Netherlands
BJ	Benin	GR	Greece	NO	Norway
BR	Brazil	HU	Hungary	PL	Poland
CA	Canada	IT	Italy	RO	Romania
CF	Central African Republic	JP	Japan	SD	Sudan
CG	Congo	KP	Democratic People's Republic of Korea	SE	Sweden
CH	Switzerland	KR	Republic of Korea	SN	Senegal
CI	Côte d'Ivoire	LI	Liechtenstein	SU	Soviet Union
CM	Cameroon	LK	Sri Lanka	TD	Chad
CS	Czechoslovakia	LU	Luxembourg	TC	Togo
DE	Germany	MC	Monaco	US	United States of America
DK	Denmark				

-1-

ARCHITECTURE FOR MINIMAL INSTRUCTION
SET COMPUTING SYSTEM

Field Of The Invention

5 The present invention relates to programmable digital computing systems. In particular, the invention relates to a digital computing system architecture that uses a minimal instruction set and minimal control circuitry to process data at significantly high rates of speed.

10 Background Of the Invention

One of the fundamental goals in any computer system architecture is to minimize the time the system takes to complete a given task. This performance measurement may be represented by the product of the 15 number of instructions required to do the task, the duration of the clock cycle, and the average number of clock cycles per instruction execution.

One traditional type of computer design strives to minimize the number of instructions required to 20 complete the task with powerful instructions. Many such computing systems implement a complex instruction set, known as Complex Instruction Set Computer (CISC) architecture. In this arrangement, the system is designed to minimize memory storage by using instructions as 25 pointers to stored control signals in the form of microprograms in read-only memory (ROM).

-2-

For example, the central processing unit (CPU) of this system may be partitioned into control and data processing units wherein the control unit provides control signals to the data processing unit for selection and 5 sequencing of data processing operations. In many such arrangements, the control unit operates in a two-step fetch-execute cycle. In the fetch step, an instruction is obtained from memory and a decoder determines the nature of the instruction. The execute step thereafter carries 10 out the necessary operations for manipulation of data through the data processing unit.

The data processing unit of these systems performs arithmetic and logical operations on data supplied thereto in response to decoded control unit 15 signals. In traditional systems, the data processing unit, or arithmetic logic unit, is designed to provide the maximum number of operations using the fewest number of gates. Accordingly, the overall propagation delay through the data processing unit is high.

20 The control unit function is provided by a microprogram wherein a sequence of operations perform specific tasks that are much more elementary than the machine instruction fetched from memory. A microprogram sequencer, the equivalent of a program counter at the 25 microprogram level, sequentially locates the address of the macroinstruction in ROM and thereafter provides the microinstructions for directly controlling functions at desired points within the CPU. The instruction code, therefore, is quite complex which results in multiple 30 clock cycles to complete the task. Accordingly, the most

-3-

sophisticated CISC systems are able to achieve a performance of executing approximately 17 million instructions per second (MIPS).

A simpler architecture has been described in 5 Reduced Instruction Set Computer (RISC) systems. Such systems utilize a smaller instruction set which have instructions designed to reduce the number of clock cycles per instruction execution as well as to reduce the instruction decode cycle. For example, many RISC systems 10 are designed to average about one clock cycle to execute an instruction. This arrangement provides simpler and faster fetch and decode hardware than found in the traditional CISC architecture.

The RISC instruction set also seeks to minimize 15 address modes and instruction formats. RISC systems use load/store architecture to support the reduction of the instruction set. Thus, only the load and store instructions, which operate to move data between registers and memory, can access memory. The simplification results 20 in an increased amount of code required to perform an operation in comparison to CISC systems. In modern microprocessors, however, computation is faster than memory access. Thus, for example, in RISC machines, load or store instructions will be started in one machine 25 cycle, and in the next cycle while the load and store is in progress, the RISC CPU can execute the next instruction.

In comparison with CISC systems, the relative simplicity of the instruction set partially eliminates the 30 need for microcode in the CPU in selected RISC systems. With use of this approach, RISC systems approach an execution performance of approximately 33 MIPS,

-4-

significantly greater compared to CISC systems. However, instruction decoding circuitry must still decode the sequences of low-level operations of the RISC system. In addition, the execution rate of such systems is limited by 5 the number of instructions supported by the processor. Thus, while such prior art systems may function satisfactorily under certain conditions, they are still quite complex.

Objects of the Invention

10 Accordingly, a general object of the present invention is to provide an improved architecture for a digital computing system that overcomes the deficiencies of the prior art.

15 Similarly, an object of the present invention is to provide a digital computing system architecture that greatly increases execution speed.

20 A further object of the present invention is to provide a computing system architecture which utilizes a minimal instruction set to implement as few instructions as possible in hardware.

25 An additional object of the present invention is to provide a computing system architecture that uses direct instruction control wherein an instruction code from memory directly stimulates control lines in the hardware to eliminate the instruction decoding step and to simplify the hardware design.

-5-

It is another object of the present invention to provide a computing system architecture that uses a data flow-through design for the hardware to minimize undesirable characteristics of a synchronously clocked 5 signal.

It is an additional object of the present invention to provide the functionality traditionally found in microcode with software to enhance user flexibility.

Other objects and advantages of the invention 10 will become apparent upon reading the following description and appended claims, and upon reference to the accompanying drawings.

Summary Of The Invention

These objects are achieved with a computing 15 system architecture that uses a small streamlined instruction set that operates at higher speeds than other known techniques. In addition, the invention applies unencoded instructions directly from memory to a processing means to provide computational paths for data 20 to flow therethrough. The method according to the present invention accomplishes this by manipulating data in a general purpose digital computing system having memory means, central processing means, and a bus network interconnecting the central processing means and memory 25 means. An instruction word having source operands of data and, destination operand of data, and an instruction portion having bit portions therein attributable to a minimum instruction set is located. Selected bits of the instruction portion control streamlined computational 30 paths in the central processing means. The bit portions

-6-

of the instruction are applied directly to the central processing means. In this way, a result may be obtained therefrom and stored in a memory location attributable to the destination operand. The invention executes simple 5 instructions at a very rapid processing speed to provide a performance advantage over conventional processing techniques. It has been found that such an arrangement may provide an execution performance of 185 MIPS.

Brief Description of the Drawings

10 For a more complete understanding of this invention, reference should now be made to the embodiment illustrated in greater detail in the accompanying drawings and described below by way of an example of the invention. In the drawings:

15 Fig. 1 is a simplified block diagram representation of the computing system architecture of the present invention implemented in a processor.

Fig. 2 is a block diagram representation of the control unit of the processor shown in Fig. 1.

20 Fig. 3 is a block diagram representation of the universal functional unit of the processor shown in Fig. 1.

Fig. 4 is a block diagram representation of the memory management unit of the processor shown in Fig. 1.

25 Fig. 5 is a single instruction pipeline flow diagram of an instruction provided by the control unit of Fig. 2.

Fig. 6 is a diagrammatic representation of a direct instruction control word in accordance with the 30 present invention.

-7-

Fig. 7 is a diagrammatic representation of a software compilation sequence to provide the processor shown in Fig. 1 with an atomic operation from the minimum instruction set used in the present invention.

5 Fig. 8 is a diagrammatic representation of a floating point number in accordance with the present invention.

Detailed Description Of The Preferred Embodiment

The following detailed description will permit a 10 more complete understanding of this invention. However, the embodiment described below is simply an example of the invention, and the invention is not limited to this embodiment. It will be understood that the architecture of the present invention may be implemented with the use 15 of various configurations of processors with appropriate modification. Accordingly, the specific configuration of the processor hereinafter described provides a nonlimiting example of one implementation of the present invention. It will be further understood, that in certain instances, 20 details may have been omitted which are not necessary for an understanding of the present invention.

Generally, the present invention relates to a high speed computing system architecture. The architecture is conceptually understood as a computing 25 machine that compiles directly from a high level language to direct instruction control words providing the functional equivalent of microcode. The present invention uses a two-layer instruction set that compiles from a high level language to atomic instructions, which contain the 30 functionality of microcode. The apparatus and method of

-8-

the present invention use this simplified atomic instruction set to reduce the complexity of control circuitry in the processor. Inasmuch as the functionality traditionally found in microcode has been transferred to 5 the atomic instruction set, the software complexity of the present invention is increased in comparison to traditional computer architectures.

One very important feature of the invention is that instruction sequencing is not required. This 10 arrangement directly results in significantly higher performance. Thus, for example, the invention eliminates macro or micro instructions as well as microprogram sequencing. In addition, the invention operates without decoding the instruction obtained from memory. The 15 absence of an intermediate decoding step as well as the absence of microprogram sequencing enhances the minimization of control circuitry. The minimal amount of control circuitry needed significantly reduces the size of the processor permitting highly increased processor 20 density. The invention does however, require an increased amount of memory in comparison to known CISC or RISC systems.

The system further uses an instruction word that contains operand data or address information, operation 25 information and control. The device of the present invention stores unencoded instruction words in memory to provide bits which directly stimulate control lines in the hardware. The instruction word provides control bits, directly from memory, that control the data flow-through 30 paths for the functional units, described in detail

-9-

herein. In addition controlling the functional units, the instruction word also contains the data operands and branching and trap control.

The architecture of the present invention is simplified to provide all of the necessary arithmetic and logical operations with functional units designed for minimal propagation delay therethrough. It will also be appreciated that the architecture contemplated by the present invention utilizes minimized access to registers.

For example, there are no general purpose registers provided in a preferred embodiment.

Turning now to the drawings, Fig. 1 shows a simplified block diagram representation of a processor 10 implementing the streamlined architecture of the present invention. The processor 10 is composed of three components preferably arranged on a single chip: a control unit (CU) 12, a universal functional unit (UFU) 14, and a memory management unit (MMU) 16. Of course, different arrangements for the processor 10 may occur to those skilled in the art to which this invention pertains. According to one preferred embodiment, the MMU 16 provides an interface between the processor 10 and user memory. The UFU 14 provides five functional units that perform the arithmetic and logical operations traditionally found in an arithmetic logic unit. The CU 12 provides an instruction pipeline that supplies the UFU 14 and provides control for timing circuitry in the processor 10.

In accordance with the present invention, a direct instruction control word 18 is fetched from memory. The instruction control word 18 of a preferred embodiment comprises 256 bits having four 64-bit segments: two source operands, shown as segments 20 and 22, a destination

-10-

operand shown as segment 24, and an instruction segment 26. The information contained in the source operand segments 20 and 22 may specify actual values upon which an operation is to be performed, the addresses of such 5 values, an offset value to be added to an address already stored in an index register in the MMU 16, or an address on a register bus (shown in Figs. 3 and 4 as numeral 77). As with the source operand segments, the destination operand segment 24 contains address information; the 10 destination operand segment, however, may not have a value stored therein.

The opcode in the instruction segment 26 is completely unencoded, so that the instruction segment 26 directly provides control bit signals to the control unit 12. The control unit 12 appropriately applies control bit 15 signals of the instruction segment 26 to control the computational paths for the source operand segments 20 and 22 and apply the same control directly to the computational elements of the universal functional unit 14 20 via a line 30. In this manner, the source operand segments 20 and 22 are applied to the UFU 14 via lines 27 and 29. When the UFU 14 has completed manipulation of the source operand segments 20 and 22, a result may be written to memory or to a register on a line 15. This arrangement 25 eliminates microcode sequencing found in traditional architectures. In those systems, an opcode represents an entry point into microcode which thereafter determines the movement of data to effect an operation.

The control unit 12 concurrently applies control 30 bits in the instruction segment 26 via line 30 to control the MMU 16 for the receipt of address information of the source operand segments 20 and 22 and the destination

-11-

operand segment 24. The MMU 16 thereby receives this address information via lines 27a, 29a and 31 and performs necessary indexing and translation operations thereon in accordance with control signals provided by the 5 instruction segment 26. The MMU 16 addresses memory via a line 17.

Turning now to Fig. 6, therein is shown a diagrammatic representation of the instruction segment 26 of direct instruction control word 18. The instruction 10 segment 26 comprises 64 bits in one preferred embodiment having a portion partitioned into an operation bit section 200 used to specify a desired operation.

One very important feature of the present invention is the use of a minimum instruction set 15 comprising seven pure atomic instructions which define the operations performed by the processor 10. The minimum instruction set allows the absence of decoding in the instruction control word 18 by reducing the word 18 to a manageable size. This is accomplished by assigning a 20 particular bit in the operation bit section 200 to one atomic instruction. The minimum instruction set also reduces the hardware complexity of the processor 10 inasmuch as the processor 10 has fewer instructions to support. These atomic instructions, shown in Table 1 25 below, enable the processor 10 to perform all the tasks of a general purpose computer:

-12-

Table 1MINIMUM INSTRUCTION SET

	<u>Instruction</u>	<u>Description</u>
5	SUM	Add (Subtract)
	PROD	Multiply
	SHIFT	Shift left or right: arithmetic, logical, or rotate; or sticky right
10	LOGIC	Logical AND, OR, NAND, NOR, XNOR, XOR
	TRAP	Trap
	UNPACK	Floating Point Unpack
	PACK	Floating Point Pack

15 It will be appreciated that the minimum instruction set performs integer arithmetic, logical, floating-point arithmetic, and control operations. Appropriate software modification permits the processor 10 to combine these atomic instructions to perform other instructions as well.

20 For example, twelve other low-level instructions may be implemented by setting bits in an option bit section 202 of the instruction segment 26. By so doing, a subtract instruction is implemented wherein the addition instruction adds the 2's complement of the data operand 25 segment 22 with the data operand segment 20. Similarly, bits in the option section 202 may be set to provide a logical shift right or left instruction, an arithmetic shift right or left instruction, a rotate left or right

-13-

instruction, and a sticky right instruction used for floating point operations. In addition, bits in the option bit section 202 may be set to provide logic AND, OR, XOR, NAND, NOR and XNOR operations.

5 Atomic operations may be locked together to form new indivisible psuedo-instructions through the use of the LOCK instruction locking prefix. The presence of a bit 212 in the instruction segment 26 corresponding to a lock instruction dictates a sequence, or a series of multiple 10 atomic instructions in a noninterruptable entity. Since locked instructions cannot be interrupted, the processor 10 will not process an interrupt until the last instruction of the locked set is finished. Thus, any other computer instruction can be made from combinations 15 of the atomic instructions.

The following table provides a description of the atomic operations performed by the present invention along with corresponding options.

-14-

Table 2

<u>INSTRUCTION</u>	<u>DESCRIPTION</u>	<u>OPTION</u>
SUM	Addition/Subtraction	Add/Subtract
PROD	Multiplication	None
5 SHIFT	Shift	Left/Right, Logical/ Arithmetic, Rotate, Sticky Right
10 LOGIC	Logical	AND, OR, NAND NOR, XNOR, XOR
TRAP	Trap	None
UNPACK	Floating Point Unpack	Sum/Product
PACK	Floating Point Pack	Sum/Product
15	In the preferred embodiment, seven bits designate the atomic operations in the operation section 200 of the instruction segment 26. The SUM, PRODUCT, LOGIC, SHIFT, UNPACK, PACK and TRAP atomic operations are assigned to bits 56 through 63, respectively, of the 20 instruction segment 26. An escape operation, which prevents the processor 10 from attempting to execute an instruction, is assigned to bit 63.	
25	Four bits designate operation options in the option bit section 202 of the instruction segment 26 and have overlapping bit functions for various operations. When the LOGIC operation bit is selected, the NOT, AND, OR, and XOR options are assigned to bits 52 through 55, respectively. When the SUM operation bit is selected, the addition or subtraction mode option is assigned to bit 52. When the 30 SHIFT operation is selected, left/right, logical/arithmetic,	

-15-

rotate and sticky right options are assigned to bits 52 through 55, respectively. When FLOATING POINT UNPACK is selected, the sum/product option is assigned to bit 52. When FLOATING POINT PACK is selected, the sum/product is also 5 assigned to bit 52.

The following is a description of the atomic instruction operations performed by the present invention wherein "Usrc" designates and unrestricted source, "Udst" designates and unrestricted destination, "Msrc" designates a 10 memory source, "Mdst" designates a memory destination, "Fsrc" designates a floating point register source (shown as numeral 82 in Fig. 3), and "Fdst" designates a floating point register destination. It will be understood that, for example, "Usrcl," "Mscr1" or "Fsrc1" refers to source 15 operand segment 20, "Usrc2," "Mscr2" or "Fsrc2" refers to source operand segment 22 and "Udst," "Mdst" or "Fdst" refers to destination operand segment 24.

1. Integer Addition

(Udst) # (Usrcl) + (Usrc2)

20 The atomic addition operation performs a two's complement addition. The data at the sources Usrcl and Usrc2 are added together and their sum overwrites the data at the destination Udst.

2. Integer Subtraction

(Udst) # (Usrcl) - (Usrc2)

25 The atomic integer subtraction operation subtracts the data at source Usrc2 from the data at source Usrcl. The data at both Usrcl and Usrc2 are assumed to be in 2's complement representation. The result overwrites the data 30 at the destination Udst.

-16-

3. Integer Multiplication

(Udst) # (Usrcl) x (Usrc2)

The atomic multiplication operation multiplies the data at source Usrcl by the data at source Usrc2. The 64 least significant bits of this result overwrite the data at the destination. Bit 65 of the result is copied to the carry flag. If any of the 63 most significant bits are set, the overflow flag is set. Otherwise, the 64 most significant bits are discarded. The data at the sources Usrcl and Usrc2 are assumed to be in 2's complement notation. The result of the multiplication is in 2's complement notation.

4. Shift Left (Logical)

(Udst) # (Usrcl) FF (Usrc2)

The atomic logical shift left operation shifts the data at source Usrcl left by the shift distance specified by the data at source Usrc2. Zeros are shifted in on the right. The result overwrites the data at destination Udst. If the data at Usrc2 is negative, an effective logical shift right occurs by the absolute value of the number of the shift distance specified by Usrc2. The shift distance used is $(Usrc2) \% 64$; i.e. only the 6 least significant bits of (Usrc2) are used for the shift distance.

5. Shift Right (Logical)

(Udst) # (Usrcl) ff (Usrc2)

The atomic logical shift right operation shifts the data at source Usrcl right by the shift distance specified by the data at source Usrc2. Zeros are shifted in on the left. The result overwrites the data at destination Udst. If the data at Usrc2 is negative, an effective logical shift left occurs by the absolute value of the number of the shift distance specified by Usrc2. The shift distance used is $(Usrc2) \% 64$; i.e. only the 6 least significant bits of (Usrc2) are used for the shift distance.

-17-

6. Shift Left (Arithmetic)

(Udst) + (Usrcl) << (Usrc2) with sign extension

The atomic arithmetic shift left operation shifts the data at source Usrcl by the shift distance specified by 5 the data at source Usrc2. Copies of the least significant bit of Usrcl are shifted in from the right. The result overwrites the data at destination Udst. If the data at Usrc2 is negative, an effective arithmetic shift right occurs by the absolute value of the number of the shift 10 distance specified by Usrc2 %64; i.e. only the 6 least significant bits of (Usrc2) are used for the shift distance.

7. Shift Right (Arithmetic)

(Udst) + (Usrcl) >> (Usrc2) with sign extension

The atomic arithmetic shift right operation shifts 15 the data at source Usrcl by the shift distance specified by the data at source Usrc2. Copies of the most significant bit of Usrcl are shifted in from the left (sign extension). The result overwrites the data at destination Udst. If the data at Usrc2 is negative, an effective arithmetic shift 20 left occurs by the absolute value of the number of the shift distance specified by Usrc2. The shift distance used is (Usrc2) %64; i.e. only the 6 least significant bits of (Usrc2) are used for the shift distance.

8. Shift Right (Rotate)

25 (Udst) + (Usrcl) rotated right (Usrc2) bits

The atomic rotate right operation rotates the data at source Usrcl right by the distance specified by the data at source Usrc2. The result overwrites the data at destination Udst. If the data at Usrc2 is negative, a 30 rotate left occurs by the absolute value of the distance

-18-

specified by UsrC2. The distance used is (UsrC2) %64; i.e. only the 6 least significant bits of (UsrC2) are used for the rotate distance.

9. Shift Left (Rotate)

5 (Udst) ← (UsrC1) rotated left (UsrC2) bits
The atomic rotate left operation rotates the data at source UsrC1 left by the distance specified by the data at source UsrC2. The result overwrites the data at destination Udst. If the data at UsrC2 is negative, a
10 rotate right occurs by the absolute value of the distance specified UsrC2. The distance used is (UsrC2) %64; i.e. only the 6 least significant bits of (UsrC2) are used for the rotate distance.

10. Shift Right Arithmetic Sticky

15 A sticky right shift also performs a right shift on the operand; if any of the discarded bits of the original are ones, the least significant bit of the result becomes a one (regardless of the value it would otherwise have). The sticky right and arithmetic right shifts can be performed
20 simultaneously on the same operand.

11. Logical AND

(Udst) ← (UsrC1) AND (UsrC2)

The data bits of the source at UsrC1 are AND'ed with the bits of the data source UsrC2. The result
25 overwrites the data at destination Udst.

12. Logical OR

(Udst) ← (UsrC1) OR (UsrC2)

The data bits of the source at UsrC1 are OR'ed with the bits of the data at source UsrC2. The result
30 overwrites the data at destination Udst.

-19-

13. Logical NAND

$(Udst) \leftarrow (Usrcl) \text{ NAND } (Usrc2)$

The data bits of the source at Usrcl are NAND'ed with the bits of the data at source Usrc2. The result 5 overwrites the data at destination Udst.

14. Logical NOR

$(Udst) \leftarrow (Usrcl) \text{ NOR } (Usrc2)$

The data bits of the source at Usrcl are NOR'ed with the bits of the data at source Usrc2. The result 10 overwrites the data at destination Udst.

15. Logical XNOR

$(Udst) \leftarrow (Usrcl) \text{ XNOR } (Usrc2)$

The data bits of the source at Usrcl are XNOR'ed with the bits of the data at source Usrc2. The result 15 overwrites the data at destination Udst.

16. Floating Point Unpack

$(Fdst) \leftarrow (Usrcl), (Usrc2)$

This atomic operation unpacks two floating point numbers at source Usrcl and Usrc2 into the destination Fdst, 20 a point floating register file 82 (shown in Fig. 3), and is described in greater detail herein.

17. Floating Point Pack

$(Udst) \leftarrow \text{Pack } (Usrcl), (Usrc2)$

This atomic operation packs an integer mantissa 25 and exponent at source Usrcl and Usrc2, into the destination Udst, and is described in greater detail herein.

18. Trap

The atomic trap instruction executes a software 30 interrupt. The contents of the program counter register are saved in a program counter shadow register in the program counter control circuit (numeral 80 in Fig. 4). The contents of the processor status register (numeral 72 in Fig. 3) are saved in a processor status shadow register (also numeral 72). The address of the common entry point

-20-

of all trap routines is placed in the program counter 80. A source operand segment 20 or 22 is placed in an interrupt vector register in the interrupt control circuit (numeral 44 of Fig. 2) to indicate the type of trap being 5 taken.

A flag enable bit section 204 is also provided in the instruction segment 26 to accommodate a variety of user modes. The flag enable bit section 204 advantageously permits a particular flag to be modified or 10 not modified on an individual instruction basis.

The flag enable designations in the flag enable bit section 204 for floating point operations: Zero (FZ), Negative (FN), Overflow (FV), Underflow (FU), Infinity (FInf), Not a Number (FNaN), Unnormal (FUnN), Inexact 15 (FInx), Invalid Operation (FInO), Divide By Zero (FDvZ) and Unordered (FUnO) are assigned to bits 28 through 38 respectively. The flag enable designations for Integer Operations: Zero (Z), Negative (N), Overflow (V), and Carry (C), are assigned to bits 24 through 27 20 respectively.

The instruction segment 26 contains a condition bit field 210 that determines whether or not the result of a computation performed by the UFU 14 should be written, based on the contents of flag bits in a processor status 25 register (shown as numeral 72 in Fig. 3). If the condition specified in the bit field 210 of instruction segment 26 is met, then the result of the computation will be written to the address specified by the destination operand segment 24. Otherwise, the result is not written. 30 This feature provides the functionality of conditional and unconditional jump instructions of a traditional computer. To perform an unconditional jump, the destination of the

-21-

5 write is a program counter register (shown in Fig. 4 as numeral 80) and the conditional bit field 210 contains a bit pattern which specifies the condition that must be met by the contents of the processor status register 72 for the write to take place. One advantage of such an arrangement is that any memory location or register can be the destination of such a conditional write operation. The use of this option is not restricted to use with only the program counter register 80 as the destination.

10 The condition designations in the condition bit field 210 may be set as follows:

Table 3

	Bit 44	43	42	41	40	39	Description
15	0	0	0	0	0	0	always
	0	0	0	0	0	1	never
	0	0	0	1	0	0	N set
	0	0	0	1	1	1	N clear
	0	0	1	0	0	0	Z set
	0	0	1	0	1	1	Z clear
20	0	0	0	1	1	0	C set
	0	0	0	1	1	1	C clear
	0	0	1	0	0	0	V set
	0	0	1	0	0	1	V clear
	0	0	1	0	1	0	GE (greater than or equal to)
	0	0	1	0	1	1	LT (less than)
25	0	0	1	1	0	0	GT (greater than)
	0	0	1	1	0	1	LE (less than or equal to)
	0	0	1	1	1	0	HI (higher)
	0	0	1	1	1	1	LOS (lower)
	0	1	0	0	0	0	FDvZ set
	0	1	0	0	0	1	FDvZ clear
30	0	1	0	0	1	0	FI _{NO} set
	0	1	0	0	1	1	FI _{NO} clear
	0	1	0	1	0	0	FI _{nx} set
	0	1	0	1	0	1	FI _{nx} clear
	0	1	0	1	1	0	FU _{uN} set
	0	1	0	1	1	1	FU _{uN} clear
35	0	1	1	0	0	0	FI _{nf} set
	0	1	1	0	0	1	FI _{nf} clear

-22-

	0	1	1	0	1	0	FNaN set
	0	1	1	0	1	1	FNaN clear
	0	1	1	1	0	0	FU set
	0	1	1	1	0	1	FU clear
5	0	1	1	1	1	0	FV set
	0	1	1	1	1	1	FV clear
	1	0	0	0	0	0	FN set
	1	0	0	0	0	1	FN clear
	1	0	0	0	1	0	FZ set
10	1	0	0	0	1	1	FZ clear
	1	0	0	1	0	0	FUnO set
	1	0	0	1	0	1	FUnO clear
	1	0	0	1	1	0	
	through						
15	1	0	1	0	0	1	unused
	1	0	1	0	1	0	FGE (floating point greater than or equal to)
	1	0	1	0	1	1	FLT (floating point less than)
20	1	0	1	1	0	1	FLE (floating point less than or equal to)
	1	0	1	1	1	0	
	through						
	1	1	1	1	1	1	unused

25 Finally, an address control section 206 of instruction segment 26 provides control for the source operand segments 20 and 22 and destination segment 24. The remaining bit section 208 is reserved.

Control in the address control section 206 for the 30 source operands 20 and 22 is the same. This control specifies if a data control bit is enabled and whether the operand is immediate or direct. If the operand is direct, whether an index is needed is determined. If an index is needed, control must select a desired index register. The 35 control also specifies if the operand is stored in memory or a register and must specify a desired segment register. Control for the destination operand 24 is identical in the preferred embodiment except that the destination operand 24 cannot be immediate; it can only be direct.

-23-

The Index/No Index, Select Index Register, Memory/Register, Select Segment Register, Immediate Direct, and Data Control Enable options for control of the data operand segment 20 are assigned to bits 16 through 5 23, respectively.

The Index/No Index, Select Index Register, Memory/Register, Select Segment Register, Immediate/Direct, and Data Control Enable options for control of the data operand segment 22 in the instruction 10 segment 26 are assigned to bits 8 through 15 respectively.

The Index/No Index, Select Index Register, Memory/Register, Select Segment Register, and Data Control Enable options for control of the destination operand segment 24 are assigned to bits 0 through 7, respectively.

15 The above assignments of the instruction segment 26 are summarized in Table 3 below:

Table 3

<u>Bit</u>	<u>Assignment</u>	<u>Bit</u>	<u>Assignment</u>
0	Index/No Index	32	FInfinity
20 1,2	Select Index Register	33	FNot a Number
3	Memory/Register	34	FUnnormal
4,5	Select Segment Register	35	FIexact
6	Data Control Enable	36	FIInvalid Operation
7	Reserved	37	FDivide by Zero
25		38	FUnordered
		39	Condition
8	Index/No Index	40	Condition
9,10	Select Index Register	41	Condition
11	Memory/Register	42	Condition
30 12,13	Select Segment Register	43	Condition

-24-

14	Immediate/Direct	44	Condition
15	Data Control Enable	45	Lock Bit
		46	Reserved
		47	Reserved
5 16	Index/No Index	48	Reserved
17,18	Select Index Register	49	Reserved
19	Memory/Register	50	Reserved
20,21	Select Segment Register	51	Reserved
10 22	Immediate/Direct	52	NOT; Add/Sub; Lft./Rt.; Sum/Prod.; Sum/Prod.
23	Data Control Enable	53	AND; Logical/ Arithmetic
15		54	OR; Rotate
		55	XOR; Sticky Right
24	Zero	56	Sum
25	Negative	57	Product
26	Overflow	58	Logic
20 27	Carry	59	Shift
28	FZero	60	Unpack
29	FNegative	61	Pack
30	FOverflow	62	Trap
31	FUnderflow	63	Escape
25	A relative jump is performed with the an instruction designating the program counter register 80 as the destination operand 24 and one of the source operands 20. The other source operand 22 is an offset to be added to the program counter register 80, and the atomic instruction performed is		
30	an addition. Absolute or computed jumps may be performed similarly.		

Atomic instructions are not normally accessed by a programmer. Instead the programmer uses an intermediate language interface, preferably the Teraplex Intermediate Language Interface (TILI), comprised of the instructions and addressing modes programmers and high level languages demand. A complete description of TILI is publicly available in TILI Reference Manual (1990) by J. Glickman from Teraplex Inc., 701 Devonshire Drive, Champaign, Illinois, and incorporated herein by reference. As shown in Fig. 7, the software architecture of the present invention takes a source language 230, such as C, and compiles it into an intermediate language such as the TILI instruction set 232. Next, the TILI program is compiled into the atomic instruction set 234. Thus, an executable program consists entirely of atomic instructions. The programmer, however, never has to write any code using atomic operations. Instead, he is free to use the TILI instruction set, or even a high-level language.

Fig. 2 shows a block diagram representation of the control means of the present invention shown as the control unit (CU) 12. When the instruction control word 18 is fetched from memory, two source operand segments 20 and 22 and the destination operand segment 24 are applied to a pipeline data/address buffer register file 32 via lines 21, 23 and 25. Similarly, the instruction segment 26 of the instruction control word 18 is applied to a pipeline instruction register file 34 via a line 36. In accordance with a preferred embodiment, control signals from the instruction segment 26 enter an execution pipeline and shift down the execution pipeline. It will be understood that the present invention would also work, albeit somewhat slower, without an execution pipeline. At each stage of the pipeline, a specific bit

portion of the instruction segment 26 provides control signals on line 30 to supply data through a desired portion of the arithmetic and logical elements of the UFU 14.

Fig. 2 also shows an instruction timing unit 50 of the CU 12. The instruction timing unit 50 is a state machine which receives control bits in the instruction segment 26 via the line 30, and an external clock via a line 46. The instruction timing unit 50 also receives asynchronous interrupt and trap signals from a trap control circuit 38 and an interrupt control circuit 42 via lines 40 and 44 as well as signals generated by the MMU 16 on a line 45. After combining the control signals with such information, the instruction timing unit 50 provides appropriate timing signals to the UFU 14 and to the MMU 16 via the line 30. The instruction timing unit also provides timing signals to other elements in the control unit 12 via a line 30a.

Another important feature of the present invention is that clocking in the processor 10 is, at least in the traditional sense, minimal. As an instruction from the instruction control word 18 begins executing, the instruction timing unit 50 enters a sequence of states that produces a clock pulse on the line 30 for a write pipeline register 78 (shown in Fig. 3) which receives the completed calculation from computation elements in the UFU 14 when the calculation of its inputs has been completed. In this manner, the flow-through cycle wherein an arithmetic, logical or control operation is performed is terminated by a synchronously clocked cycle. Accordingly, each instruction executes in minimal time, regardless of the amount of execution time required by other instructions.

-27-

The provision of clocking the system in this manner reduces the amount of wasted time while signals settle prior to a synchronous clock. In addition, this arrangement reduces the number of clock relative skews which are problematic in 5 prior art systems.

Fig. 2 also shows an illegal instruction detection circuit 35 of control unit 12. The illegal instruction detection circuit 35 directly receives control signals from the instruction segment 26 located in the pipeline instruction 10 register file 34 on the line 30 and, upon the detection of invalid instruction signals, provides an output signal to a trap control circuit 38 on a line 37 indicative of the detection of an invalid instruction.

The trap control circuit 38 receives a control bit 15 signal from the instruction segment 26 on the line 30 corresponding to execution of the atomic trap instruction. The trap control circuit 38 also receives the output signal from the illegal instruction detection circuit 35 on the line 37; and index and segment bounds check information on data in 20 the instruction control word 18 from the MMU 16 via a line 43. The trap control circuit thereafter operates to provide a signal on the line 40 to the instruction timing unit 50 indicative of an error trap or trap instruction.

An interrupt control circuit 42 is also provided in 25 the control unit 12 to process interrupts to the processor 10. An interrupt is an event that occurs asynchronously to the execution of the processor 10. Typically, peripheral devices provide an interrupt request signal via a line 47, representing a significant change in the state of the 30 peripheral device. The interrupt control circuit 42 also receives data from the two source segments 20 and 22 and destination segment 24 on lines 21, 23 and 25 as well as

-28-

control signals from the instruction segment 26 on line 30. The interrupt control circuit 42 thereafter provides an output signal on line 44 to the instruction timing unit 50 and an interrupt acknowledge signal on a line 49.

5 It will be understood by those skilled in the art to which this invention pertains that reset circuitry may also be provided to reset the elements in the control unit 12 for initializing the processor 10.

Turning now to Fig. 3 and also to Fig. 1, therein is
10 shown a block diagram representation of the functional unit means shown as the universal functional unit (UFU) 14. One important feature of the present invention is that the circuitry of the UFU 14 may be greatly simplified in comparison to the arithmetic logic units of prior art systems
15 since the UFU 14 supports the streamlined instruction set. The UFU 14 receives the two source operand segments of data 20 and 22 of the instruction control word 18 via lines 27 and 29 in a pipeline fashion and thereafter provides the functionality found in the arithmetic logic unit of
20 traditional computers. The computational path of the data received by the UFU 14 on lines 27 and 29 is determined by bits in the instruction segment 26 applied to the control unit 12 which controls such paths in the UFU 14 via the line 30.

The UFU 14 contains arithmetic elements including an
25 adder/subtractor circuit 60, a multiplier circuit 62, logical and shift elements shown as a logic/shift circuit 64, and a floating point unpacker 68, and floating point packer circuit 66 to perform floating point operations. The result of any arithmetic or logical operation of data received by the UFU 14
30 is latched to a write pipeline register 78 via a line 70. The write pipeline register 78 also receives control signals from

-29-

the control unit 12 via line 30. Additionally, the UFU 14 contains a floating point register file 82 used to store intermediate results during floating point operations.

The UFU 14 also contains the processor status 5 register 72 which provides information on the processor 10. The processor status register 72 receives status flag information from the computational elements in the UFU 14 on a line 74 and provides such information to the program counter 80 (shown in Fig. 4) as will be understood by those skilled in 10 the art. The processor status register 72 communicates with other registers in the processor 10 via a register bus 77.

Inasmuch as all of the functional units are comprised of combinatorial circuitry, an arithmetic or logical operation will be performed concurrently with the receipt of 15 data to the respective input lines 27 and 29. Accordingly, no clocking is applied to the functional units. Thus, for example, a result appears at the output line 70 of the adder/subtractor circuit 60 before a result at the output line 70 of the multiplier circuit 62. The propagation delay for 20 any computational path is strictly a function of the gate delays and other intrinsic delays of the functional unit associated therewith.

The adder/subtractor circuit 60 accepts two 64 bit binary inputs on lines 27 and 29 and produces a 64 bit output 25 at line 70. If subtraction is desired, then a subtraction mode bit is supplied by the instruction segment 26 to the adder/subtractor circuit 60 on line 30. One of the 64 bit data inputs received on line 27 is thereby complemented and then added to the other data input received on line 29.

30 The multiplier circuit 62 of the preferred embodiment is fully described in the copending application having serial number _____. As with the adder/subtractor

-30-

circuit 60, the multiplier circuit 62 accepts two binary two's complement 64 bit data inputs from the source operand segments 20 and 22 via lines 27 and 29 and produces a 64 bit output at line 70 consisting of the least significant portion of the 5 product. The multiplier circuit 62 does not manipulate carry bits from the partial product matrix during the course of addition to increase speed. Rather, the multiplier circuit 62 saves all carry bits until the final stage, where it adds the sum and carry word. While integer operations are performed 10 directly by the adder or subtracter 60, multiplier 62 and logic/shift circuits 64, integer division is performed with a series of atomic operations.

The logic/shift circuit 64 receives two 64 bit binary inputs via lines 27 and 29 and produces a 64 bit output 15 on line 70. The logic/shift circuit 64 performs six logic functions upon receipt of the appropriate control from the instruction segment 26: AND, OR, XOR, NAND, NOR and XNOR. The logic/shift circuit 64 also performs left and right arithmetic shifts, logical shifts, and circular shifts and a sticky right 20 shift. Logic functions require two 64 bit inputs at lines 27 and 29 and produce single 64 bit result at line 70. The function desired is enabled by bits in the instruction control word 18. The shift operations require one 64 bit input and the shift amount and produce a 64 bit result.

25 The preferred embodiment of the present invention supports ANSI/IEEE 754 floating point computations with the use of circuitry shown as functional blocks 68 and 66 to provide a sequence of unpack and pack operations. Floating point computations take place completely on the main chip of 30 processor 10 to eliminate an external floating point coprocessor. However, floating point computational speed could be increased with the addition of a floating point

-31-

coprocessor. Two atomic instructions, floating point unpack (UNPACK) and floating point pack (PACK), make this possible. The floating point unpack instruction takes two floating point numbers and separates each of the floating point numbers into 5 an integer mantissa and an integer exponent which are stored in the floating point register file 82. This unpack operation may be performed in one of two modes: unpack for floating point addition or unpack for floating point multiplication.

These integer values can be manipulated by the 10 adder/subtractor 60, multiplier 62, or logic/shift circuits 64. After all the operations have been performed on the mantissas and exponents, the floating point packer circuit 66 combines the mantissa and exponent to form a new floating point word. Accordingly, a floating point operation will 15 require the sequence of an unpack instruction, a sequence of integer instructions, and a pack instruction.

The floating point unpacker circuit 68 performs the first operation for floating point multiplication, addition, or subtraction. The unpacker circuit 68 accepts two floating 20 point operands at lines 27 and 29 and produces integer quantities that are manipulated with the adder/subtractor circuit 60, multiplier circuit 62, and logic/shift circuits 64, as described above. A control bit from the ICW 18 provides the unpacker circuit 68 with an unpack instruction 25 via control line 30a for either addition or multiplication.

Unpacking for addition consists of producing two signed mantissas, a biased exponent, and a shift count. The two signed mantissas are produced by merging the most significant bit and the 52 least significant bits of the 30 original 64 bit word for each word.

Data representing a floating point number, as shown in FIG. 9, contains one sign bit 216, eleven bits for the exponent 218 and 52 bits for the mantissa 220. A floating point unpack operation for addition simply separates the 5 mantissa from the exponent. The largest possible result from an addition of the two mantissas would be a 53 bit number: 52 bits in a new mantissa and a carry bit. The software of the present invention acknowledges the carry bit as fits its particular applications, e.g., by adjusting the exponent, 10 shifting the mantissa to the right and setting the most significant bit; or by generating an overflow error.

Unpacking for multiplication produces four signed mantissa sections and two biased exponent values. For multiplication operations, the product of two 52-bit integers 15 could be represented by as many as 104 bits. Therefore, the floating point unpacker circuit 68, upon receipt of a multiplication instruction, separates each mantissa into two 26-bit sections. The processor 10 performs the following four different multiplication calculations. First, the multiplier 20 circuit 62 multiplies the least significant sections of the original mantissas. Next, it multiplies the most significant sections. Finally, it multiplies the most significant section by the least significant section for each operand. After shifting and adding operations, the result is a valid product 25 wherein no bits are lost. If a fast multiplication is desired, then only two of the four 26-bit mantissa words are multiplied together and the exponents added.

Following the unpack operation, an atomic logical shift operation aligns the mantissas, if necessary. Then, as 30 acquired, integer addition and multiplication operations are performed on the mantissas and exponents. The results of these integer manipulations are placed in the write pipeline

-33-

register 78 via the line 70 and thereafter received on the register bus 77 for placement into registers in the floating point register file 82. The floating point packer circuit 66 thereafter receives the calculated mantissa and exponent words 5 from the register bus 77 (which may access lines 27 and 29) to obtain a floating point result, and then produces a reconstructed floating point number on the line 70.

Preferably, a special case detector exists in the floating point unpacker circuit 68 to determine if a floating 10 point number is not a number, infinite, or a zero. If any of these conditions occur, the appropriate floating point flag is set in the processor status register 72 via the line 74.

The bits in the flag portions of the processor status register 72 are set or cleared based on signals 15 provided by the functional blocks 60, 62, 64, 66 and 68 of the UFU 14 on the line 74. The information provided on line 74 is indicative of the result of calculations in the UFU 14. Integer operations, for example, affect the integer flag bits: 0, negative, carry and overflow. Floating point pack and 20 unpack operations affect the floating point flag bits: zero, negative, overflow, underflow, infinity, not a number, a normal inexact, invalid operations, unordered and division by zero. Flags may be disabled by resetting the flag enable bits in the instruction control word 18. This will leave the flags 25 unaffected by operations performed by the UFU 14.

Fig. 4 shows the memory management means in accordance with the present invention, shown as the memory management unit (MMU) 16. The MMU 16 provides an interface between the processor 10 and user memory. The MMU 16 30 comprises a program counter register 80, and index registers

84. The MMU 16 also comprises an index adder circuit 90, segment registers 92, a bounds verification circuit 86 and a virtual to physical address translation circuit 88.

Another feature of the present invention is that the 5 processor 10 does not have general purpose registers. The present invention contemplates use of a high speed cache memory architecture, to achieve the same performance with memory references that traditional computers achieve with registers. Accordingly, frequently accessed data items are in 10 the cache and can be retrieved from there as quickly as they can from an internal register. It will be understood, of course, that many different memory architectures may be utilized with the present invention.

Performance of the present invention, however, 15 depends on its entire support system. For example, if the processor has a very rapid processing speed, but the memory system causes wait states due to access-time restrictions, processor performance is meaningless. Therefore, the present invention involves integrating a very high speed computing 20 machine into a high speed memory system architecture.

The processor 10, however, does have special purpose internal registers directly accessible to the user: the program counter 80, index registers 84, segmentation registers 92, the floating point register file 82 (shown in Fig. 3), and 25 the processor status register 72 (also shown in Fig. 3). These registers are directly accessible to the user for reading or writing over a register bus shown as line 77.

The program counter register 80 contains the address of the next instruction to be fetched from memory into the 30 instruction pipeline. The next value of the program counter 80 is determined by the MMU 16 based on contents of the instruction control word 18 provided by the line 30 and the

-35-

contents of the processor status register 72. In the preferred embodiment, the program counter 80 is automatically incremented after each instruction fetch, unless a jump, trap or interrupt instruction is received.

5 A dedicated address provides a common entry point for all trap instructions, and a separate dedicated address for all interrupt instructions. When a trap or interrupt instruction is received, the appropriate address is automatically placed in the program counter 80 so that
10 execution begins the any correct entry point. The contents of the program counter 80 are forced to zero upon reset of the processor 10.

15 The index registers 84 hold values that are used in indexed addressing mode. This addressing mode calculates a virtual address by adding the contents of an index register and the contents of the operand's address field in a 256 bit instruction word. The index registers can be read and written over the register bus 77.

20 Preferably, one of the index registers is a Z index register used to add zero to the contents of the operand address field if an index addition from the index registers is not needed. This is a read-only register.

25 A stack pointer is another index register. It is used to implement a last-in first-out (LIFO) stack in memory as will be appreciated by those skilled in the art.

In addition, a frame pointer register is a spare index register that may be used by a compiler such as TILI as a frame pointer to ease stack frame references.

30 The processor 10 of the preferred embodiment has memory segmentation to restrict the virtual address base that programs may access. The segmentation scheme differs from that used in most traditional computers, which usually require

adding the contents of a segment register to the address. The method used in the processor 10 is much faster because this addition is not required. The address comparison takes place in parallel with paging, so that they effectively are
5 accomplished in zero time.

The processor 10 also has four pairs of segment registers 92 that are used for virtual address segmentation. These segment registers 92 can be read and written over the register bus 77.

10 The index adder circuit 90 also receives control from the control unit 12 on the line 30 to perform appropriate indexing operations on the addresses of source operand segments 20 and 22 and the destination operand segment 24, received on lines 27a, 29a, and 31 respectively. The index
15 adder 90 also performs operations on data received from the index register file 84 on the line 85. The index adder thereafter produces an output on a line 87 for bounds verification.

20 The bounds verification circuit 86 receives an output from the segment register file 92 on a line 93 specifying the virtual addressing bounds. Such information is based on control bits in the operand address control section 206 of instruction segment 26 provided on the line 30. If such a comparison produces an illegal memory reference, an
25 appropriate output signal is provided to the trap control circuit 38 on the line 87.

Otherwise, a signal provided by the bounds verification circuit 86 on a line 89 is received by the virtual to physical address translation circuit 88. As with
30 other elements in the MMU 16, the translation circuit receives control on the line 30 provided by control bits in the

-37-

instruction section 26. The address translation circuit thereafter addresses the physical location in memory on a line 17.

Paging is the process of converting virtual
5 addresses to physical addresses. The MMU 16 performs paging
with the use of virtual address as a key into a page table
that is implemented with a content addressable memory. If the
key is matched, the associated data is the physical address.
Otherwise, a page fault trap signal appears on line 43, and
10 the operating system must update the page table to contain an
entry for the new virtual address. For greatest efficiency,
two identical page tables are preferred. This eliminates
contention for page tables.

Fig. 5 shows one important feature of the CU 12,
15 providing an execution pipeline to the UFU 14. It will be
appreciated that use of the pipeline hereinafter described is
only one implementation of the present invention inasmuch as
the processor 10 will function somewhat slower without use of
a pipeline. In addition, other methods for constructing an
20 execution pipeline will occur to those skilled in the art to
which this invention pertains. Accordingly, in a single
pipeline instruction flow in accordance with one embodiment,
each instruction passes through seven stages: a page table
lookup for instruction stage 100, an instruction fetch from
25 memory stage 102, an index stage 104, a page table lookup for
operand addresses and destination index stage 106, an operand
fetch from memory and data lookup stage 108, a calculation
stage 110, and a result write to memory stage 112.

During the page table lookup for instruction fetch
30 stage 100, the execution address in the program counter
register 80 is looked up in the page table. When in user
mode, the execution address is compared against contents of

-38-

segmentation registers which contain boundary values. At the completion of this first lookup stage 100, a physical address is ready to be used for a long word mode memory fetch.

During the instruction fetch from memory stage 102, 5 a long word mode memory access fetches the next instruction control word 18 (shown in Fig 7) from memory. The instruction control word 18 consists of four 64 bit pieces: the instruction opcode segment 26, the destination operand 24, and the two source operand segments 20 and 22. The source operand 10 segments 20 and 22 may be immediate, direct, index, or register.

During the index stage 104, indexing is performed in parallel on the source operand segments 20 and 22. Bits in the instruction control word 18 specify whether the addressing 15 mode is direct or indexed. In the latter case, additional bits specify which index register is to be used. If the addressing mode is direct, indexing is performed by adding zero. At the completion of the index stage 104, virtual addresses are available for the operands 20 and 22.

20 During the data lookup 106 for operand addresses, the virtual addresses of the operand segments 20 and 22 are looked up in the page tables. If a two page table design is used, this can be done in parallel. In user mode, each virtual address is also compared with the bounds in the 25 segment register set specified by bits in the instruction control word 18. This comparison is done in parallel with the page table lookups.

Additionally, the destination address is indexed during this stage. Bits in the ICW 18 specify the desired 30 index register. At the completion of this stage 106, physical

-39-

addresses are available for the source operand segments 20 and 22 and the destination operand virtual address 24 has been completed.

During the operand fetch from memory stage 108, the 5 source operand segments 20 and 22 are fetched from memory in short word mode. The amount of time required for this operation is determined by the amount of contention for memory banks that exist between operands. At the conclusion of this stage, both source operands 20 and 22 are available to the UFU 10 14 for manipulation thereon.

In this stage, the destination virtual address is also looked up in the page table, and compared with segment register set information provided by the ICW 18. In this manner, the destination physical address is read to be used 15 for write back.

During the calculation stage 110, the calculation specified by the instruction segment 26 of the instruction control word 18 is performed. When the calculation is completed, the result is latched into the write pipeline 20 register 78.

During the result write to memory stage 112, the contents of the write pipeline register 78 are written to memory. The page and table lookup and segment bounds check for the destination address were performed in the page table 25 lookup stage 108a, and the physical address has been saved for use in this stage 112.

If one of the source operands 20 or 22 is immediate, the value bypasses the index stage 104, the page table lookup stage 106 and the read stage 108 and is held in a register 30 until it is needed in the calculation stage 110. If one of

-40-

the source operands is a processor register, the contents of that register are provided to the UFU 14 at the beginning of the calculation stage 110.

It is contemplated that the present invention will 5 provide a multiple instruction pipeline flow which is deemed to be within the skill of the art to which this invention pertains. According to the present invention, the instruction pipeline allows five separate instructions to be in various stages of completion at a given time. Such a multiple 10 instruction pipeline flow enhances the high performance of this invention. The processor instruction pipeline consists of two phases.

A novel computing system architecture meeting the aforesaid objects has therefore been described. The 15 invention provides for a minimal instruction set which results in reduced hardware complexity. Unencoded instructions are fetched from memory in a very long instruction word to provide bits which directly stimulate control lines in the hardware. It will be understood by those skilled in the art to which 20 this invention pertains that various modifications may be made in the details and arrangements of the processes and of the structures described herein in order to explain the nature of the invention without departing from the principles of the foregoing teachings. Accordingly, the invention will only be 25 limited as expressed in the appended claims.

What is claimed is:

-41-

1. A method for manipulating data in a general purpose digital computing system having memory means, central processing means, and a bus network interconnecting said central processing means and said memory means, comprising the steps of:

providing a minimal executable instruction set for functioning said digital computing system;

locating an instruction word in said memory means having first and second source operand portions, a destination operand portion and an instruction portion having preselected bit portions therein attributable to said minimal instruction set;

assigning said preselected bit portions for controlling computational paths in said central processing means;

applying ones of said preselected bit portions to said central processing means for directing data attributable to said source operand portions through a selected one of said computational paths to provide a result therefrom; and

storing said result in a location attributable to said destination operand portion.

2. The method of claim 1 wherein said step of applying further comprises the step of:

synchronously latching said result into write register means upon completion of said result.

3. The method of claim 2 wherein said source operands of data are directed through said preselected computational path independent of synchronous clocking means.

-42-

4. An instruction word having bit portions for manipulating data in a general purpose digital computing system with memory means for storing the instruction word, central processing means having control lines for controlling computational paths therein, and a bus network interconnecting said central processing means and said memory means, the word comprising:

first and second source operand portions having information attributable to data or an address where such data resides;

a destination operand portion having information attributable to the location of a destination address; and

an instruction having bit portions directly stimulating said control lines for controlling said source operand portions through preselected ones of said computational paths in said central processing means and for providing a result to said destination address.

5. A digital computing system for implementing a minimum instruction set having main memory, central processing means, and a bus network interconnecting said memory and said central processing means, said central processing means comprising:

memory management means for locating an instruction word having first and second source operands of data, a destination operand of data and an instruction having bit portions attributable to said minimum instruction set in said memory and supplying data to said memory;

functional unit means having a plurality of computational elements for performing arithmetic and logic operations on source operands of data supplied thereto; and

-43-

control means directly responsive to said instruction bit portions for providing control signals attributable to said bit portions to said functional unit means and to said memory management means.

6. The digital computing system of claim 4 wherein said functional unit means comprises:

adder/subtracter circuit means responsive to said instruction bit portions for receiving said two source operands of data and for performing addition or subtraction operations thereon and providing an output;

multiplier circuit means responsive to said instruction bit portions for receiving said two source operands of data and for performing multiplication operations thereon and providing an output;

logic/shift circuit means responsive to said instruction bit portions for receiving said two source operands of data and for performing logical and shifting operations thereon and providing an output; and

floating point computational circuit means responsive to said instruction bit portions for receiving said two source operands of data and for performing floating point operations thereon and providing an output.

7. The digital computing system of claim 5 wherein said functional unit means comprises combinatorial circuitry.

8. The digital computing system of claim 6 wherein said functional unit means further comprises write register means responsive to a synchronous signal provided by said control unit means and coupled with said adder/subtracter circuit output, said multiplier circuit output, said

-44-

logical/shifter circuit output and said floating point computational circuit output for storing the computational results therefrom.

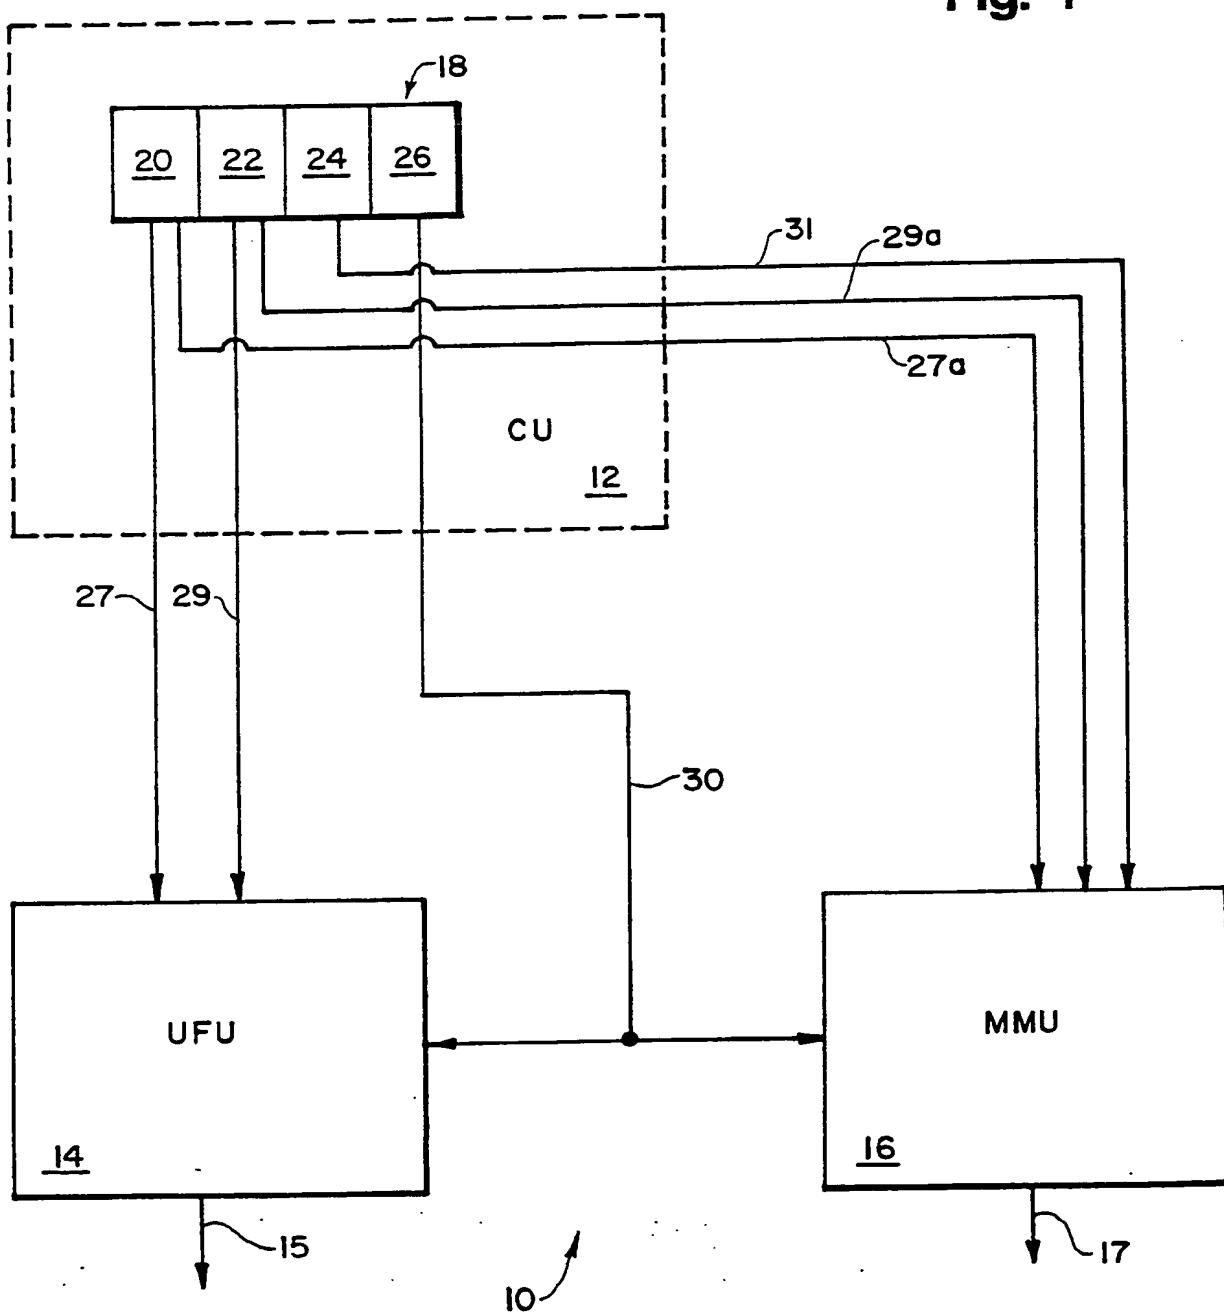
9. The digital computing system of claim 6 wherein said floating point computational circuit means comprises:

unpack circuit means responsive to said instruction bit portions for separating said source operands of data indicative of floating point data into integer exponents and integer mantissas; and

pack circuit means responsive to said instruction bit portions for forming a floating point result from said integer exponents and integer mantissas upon manipulation thereon by said adder/subtracter circuit means, said multiplier circuit means, and said logic/shift circuit means.

1/6

Fig. 1



2/6

Fig. 2

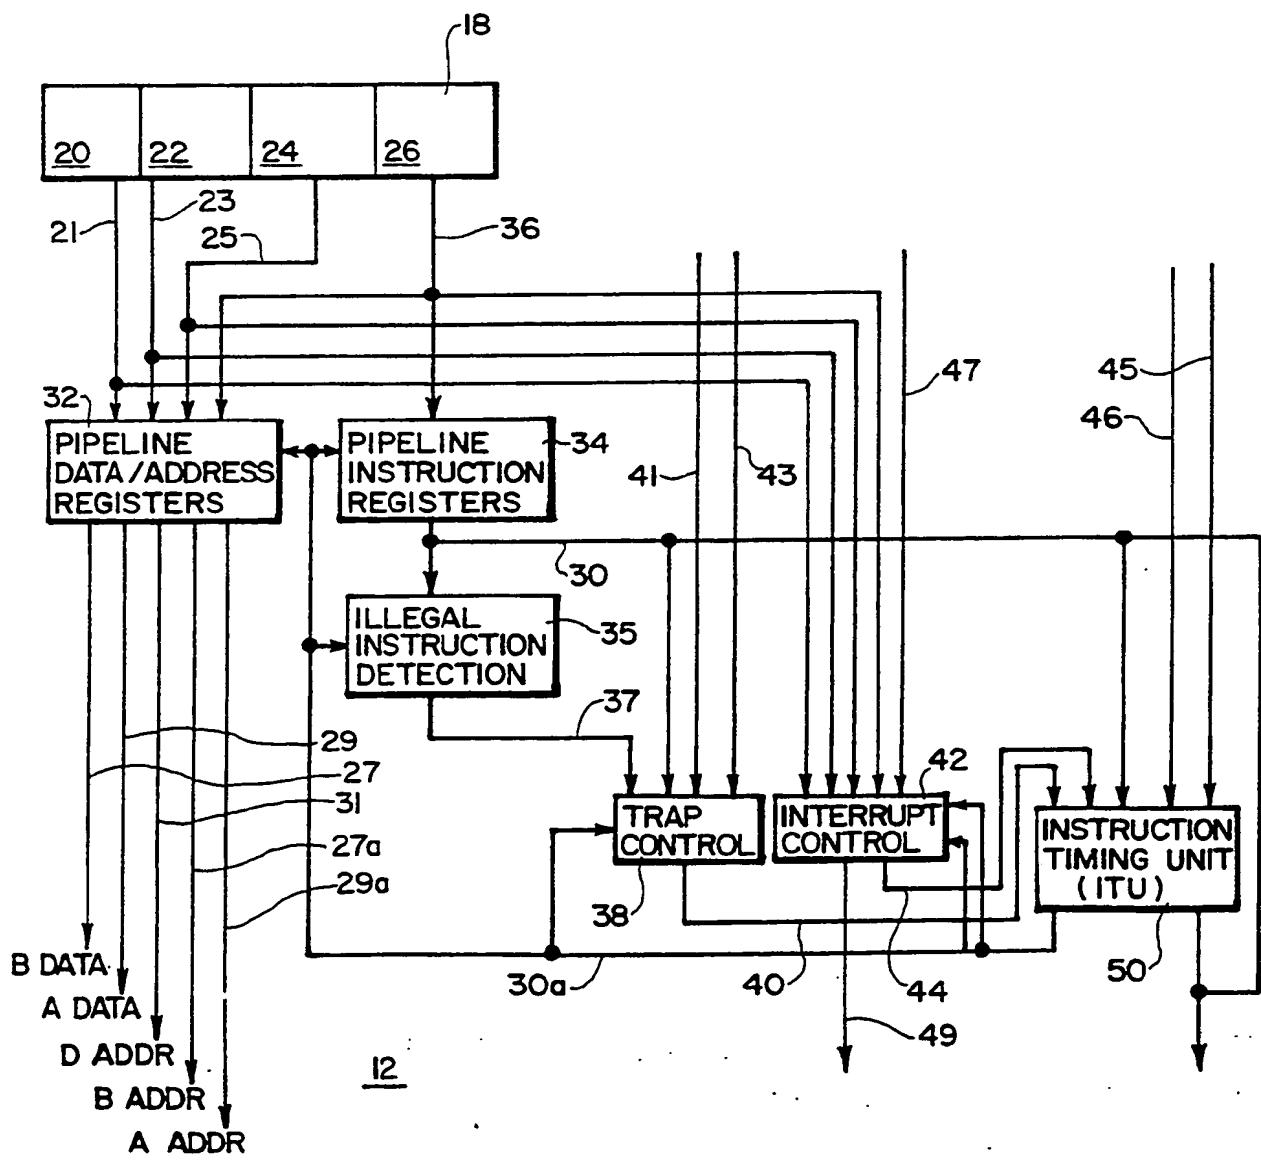
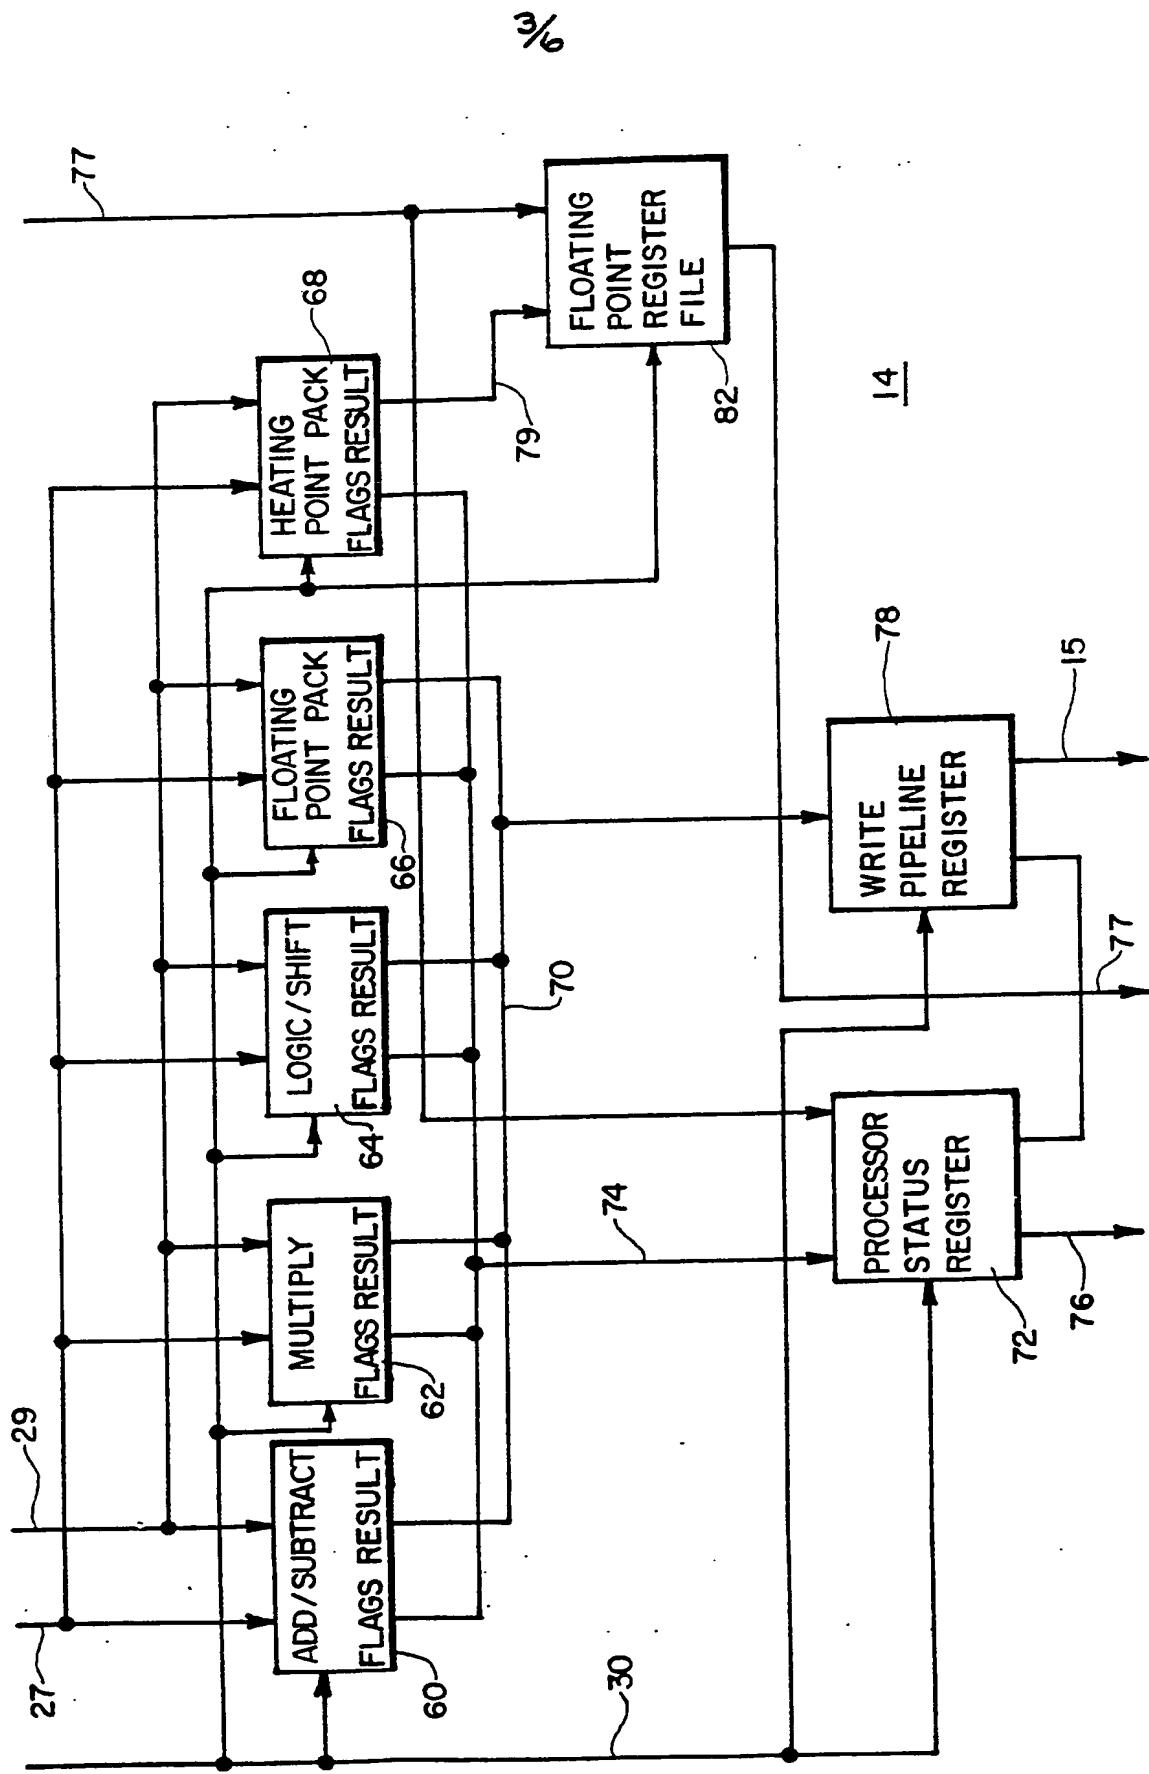
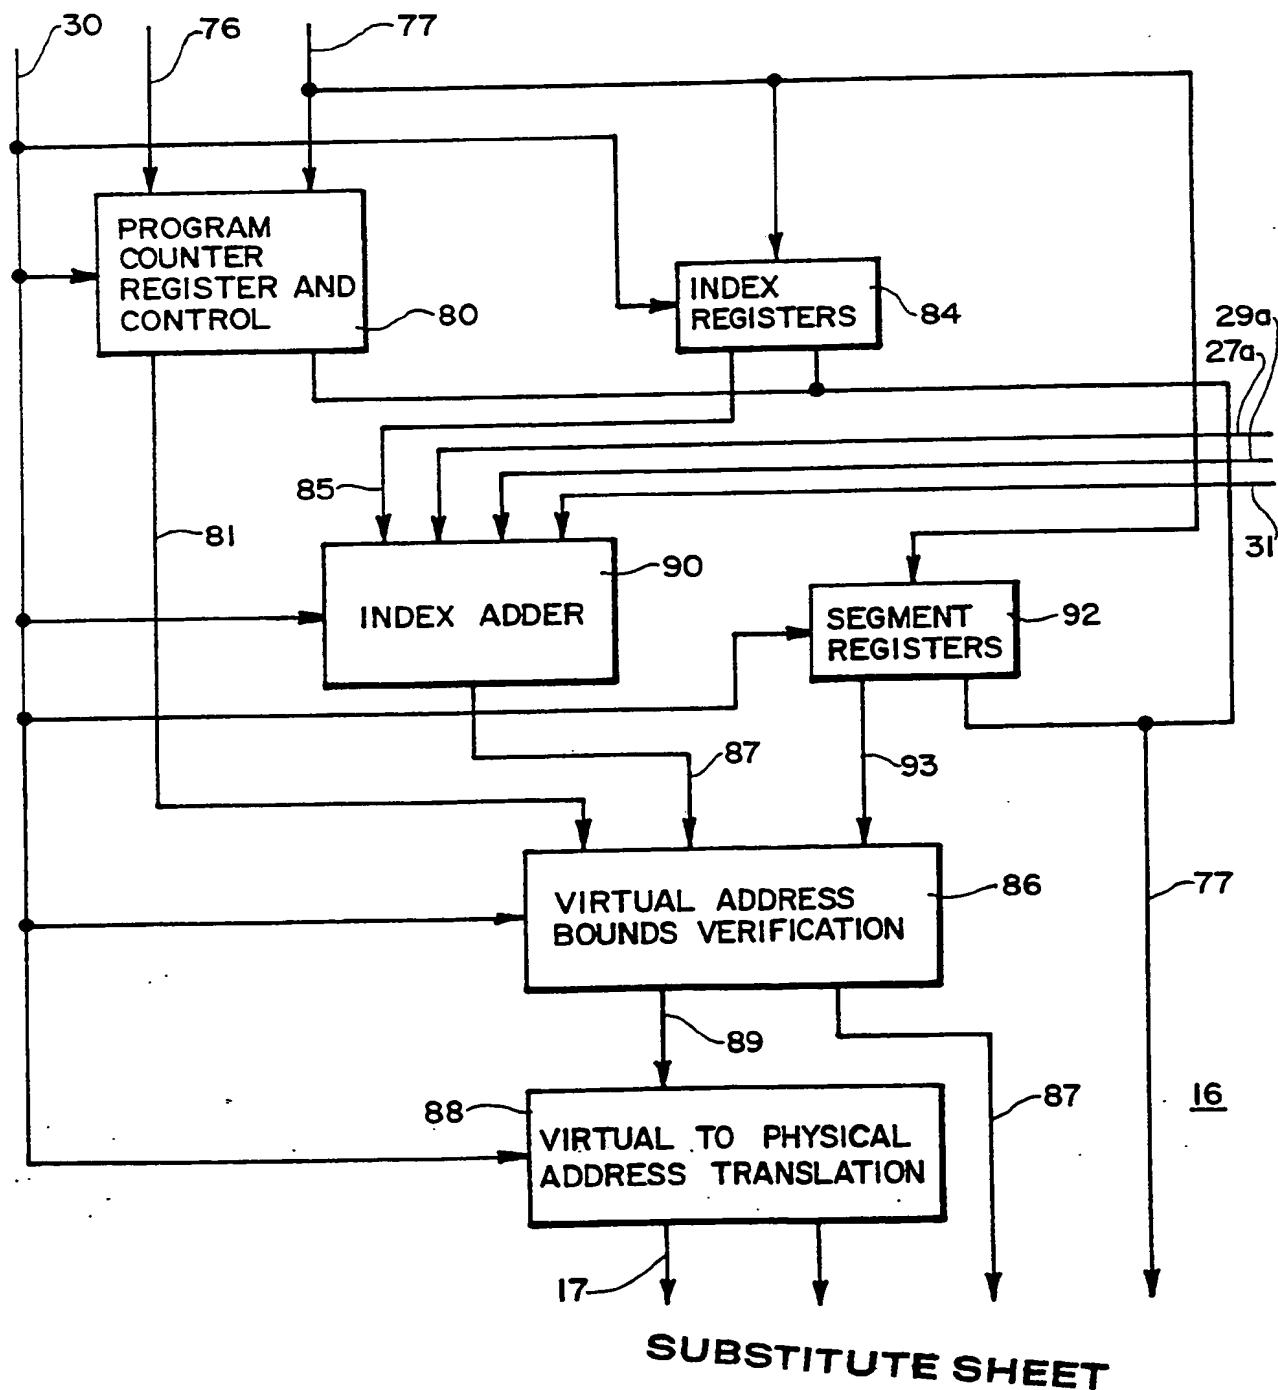


Fig. 3



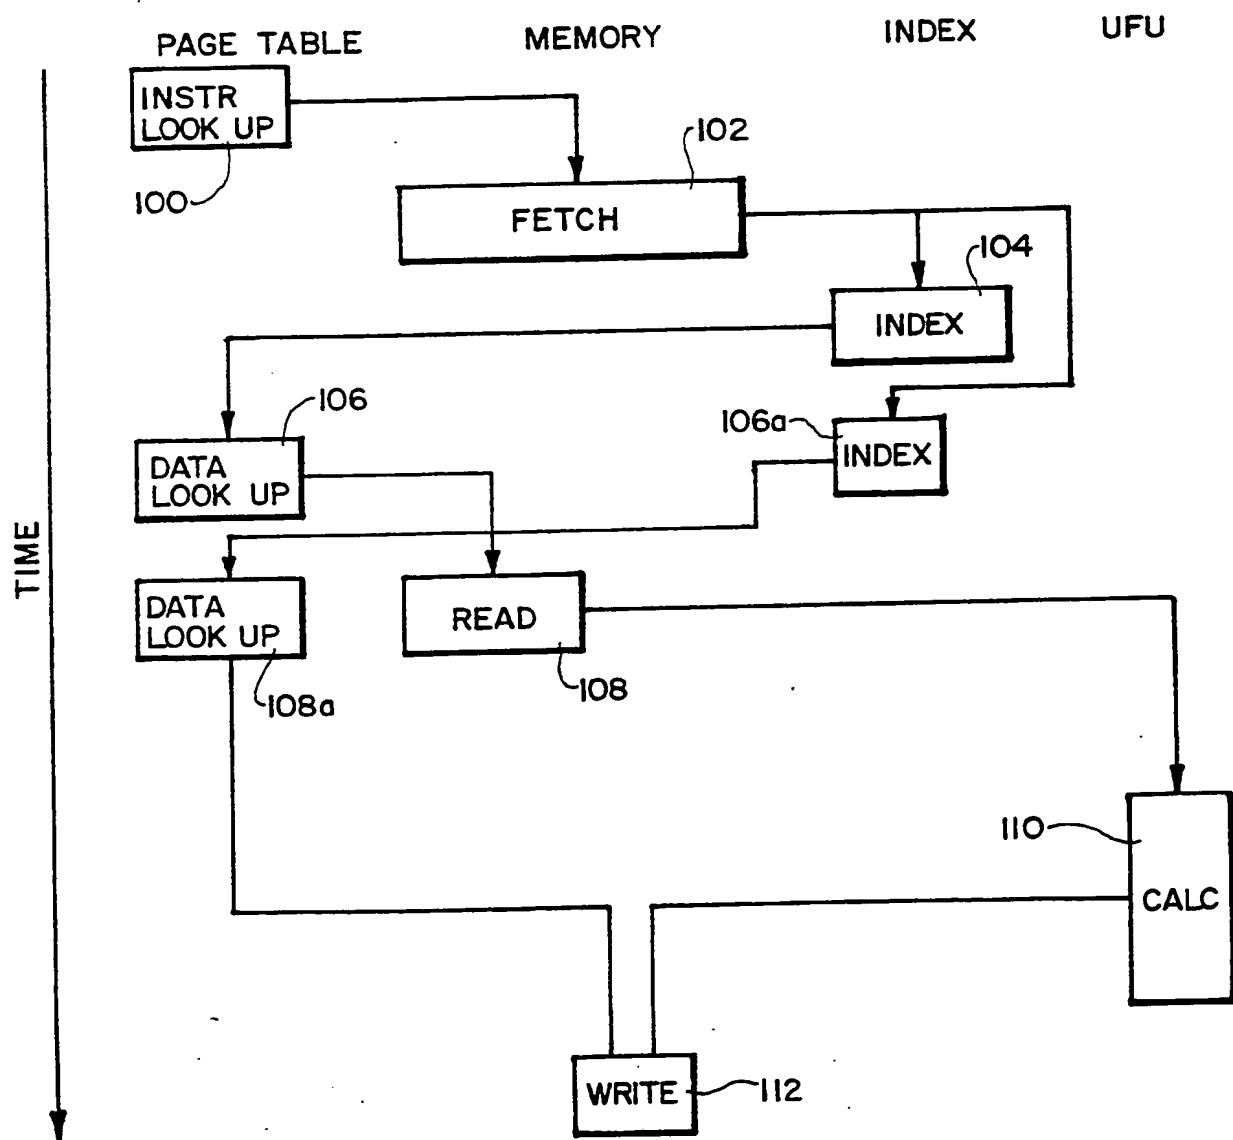
4%

Fig. 4



5/6

Fig. 5



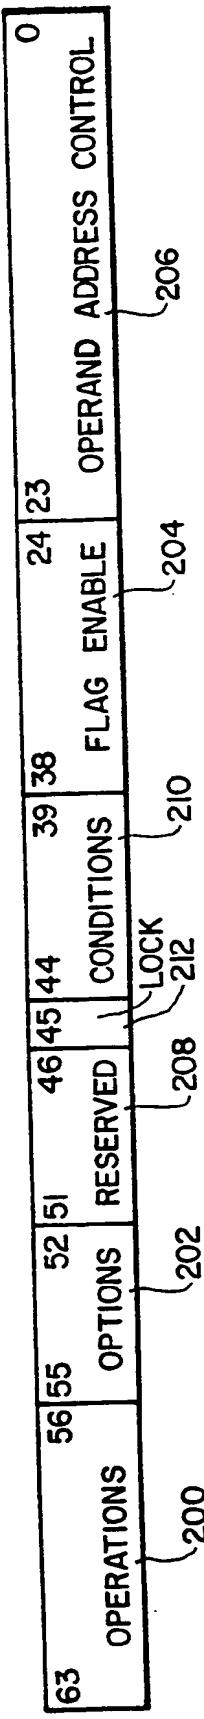


Fig. 6

6/6

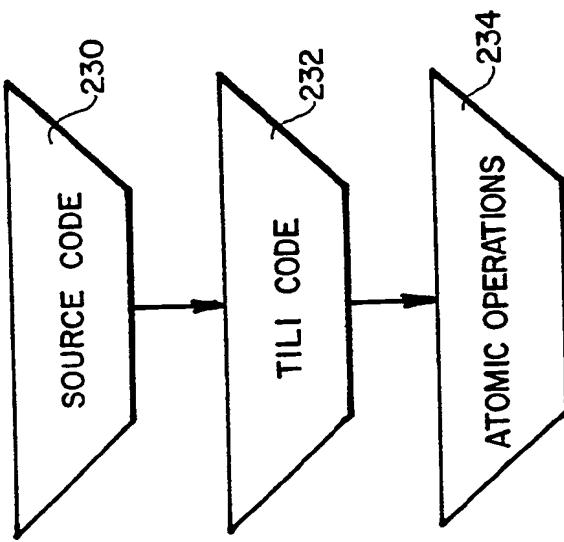


Fig. 7

SUBSTITUTE SHEET

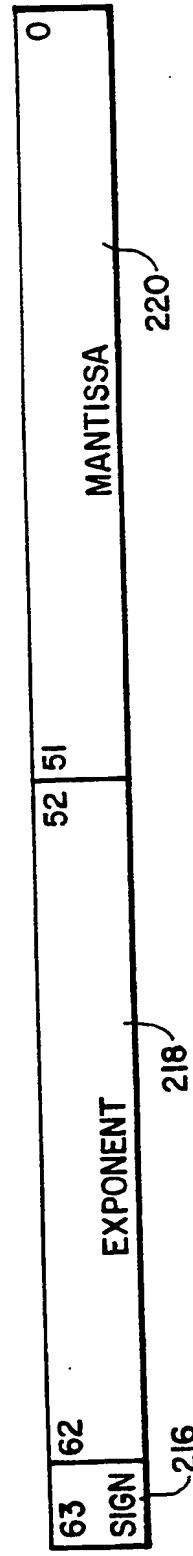


Fig. 8

INTERNATIONAL SEARCH REPORT

International Application No PCT/US91/00587

I. CLASSIFICATION OF SUBJECT MATTER (if several classification symbols apply, indicate all) ¹³

According to International Patent Classification (IPC) or to both National Classification and IPC

IPC(5): G06F 9/40
U.S. CL.: 364/200

II. FIELDS SEARCHED

Minimum Documentation Searched ¹⁴

Classification System	Classification Symbols
U.S.	364/200, 900 "MS FILE"
Documentation Searched other than Minimum Documentation to the Extent that such Documents are Included in the Fields Searched ¹⁵	

III. DOCUMENTS CONSIDERED TO BE RELEVANT ¹⁶

Category ¹⁷	Citation of Document, ¹⁸ with indication, where appropriate, of the relevant passages ¹⁹	Relevant to Claim No. ¹⁹
Y A	U.S. A, 4,734,852 (JOHNSON ET AL.) 29 March 1988, See the abstract, col. 4, lines 9-25, 30-51, col. 5, lines 56-65, col. 6, lines 33-45.	1-5, 7, AND 8 6 AND 9
Y A	U.S. A, 4,594,659 (GUENTHNER ET AL.) 10 June 1986, See col. 6, lines 36-55.	1,4,5, 2,3,6-8
X,P	U.S. A, 4,926,323 (BAROR ET AL.) 15 May 1990, See the entire document.	1-9
A	U.S. A, 4,851,990 (JOHNSON ET AL.) 25 July 1989, See the entire document.	1-9
A	U.S. A, 4,739,471 (BAUM ET AL.) 19 April 1988, See the entire document.	1-9

* Special categories of cited documents: ¹⁵

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier document but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

IV. CERTIFICATION

Date of the Actual Completion of the International Search ²⁰

20 APRIL 1991

Date of Mailing of this International Search Report ²¹

13 MAY 1991

International Searching Authority ²²

ISA/US

Signature of Authorized Officer ²³

AYNI MOHAMED

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- BLACK BORDERS**
- IMAGE CUT OFF AT TOP, BOTTOM OR SIDES**
- FADED TEXT OR DRAWING**
- BLURRED OR ILLEGIBLE TEXT OR DRAWING**
- SKEWED/SLANTED IMAGES**
- COLOR OR BLACK AND WHITE PHOTOGRAPHS**
- GRAY SCALE DOCUMENTS**
- LINES OR MARKS ON ORIGINAL DOCUMENT**
- REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY**
- OTHER:** _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.